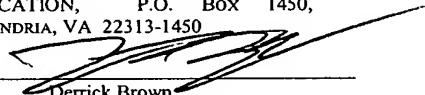


**PATENT
5681-70200
P9262**

"EXPRESS MAIL" MAILING LABEL
NUMBER EV317117063US
DATE 8-19-03 OF DEPOSIT

I HEREBY CERTIFY THAT THIS PAPER
OR FEE IS BEING DEPOSITED WITH
THE UNITED STATES POSTAL
SERVICE "EXPRESS MAIL POST
OFFICE TO ADDRESSEE" SERVICE
UNDER 37 C.F.R. §1.10 ON THE DATE
INDICATED ABOVE AND IS
ADDRESSED TO THE COMMISSIONER
FOR PATENTS, BOX PATENT
APPLICATION, P.O. Box 1450,
ALEXANDRIA, VA 22313-1450


Derrick Brown

COMPUTER ASSEMBLY

By:

James Edward King
Rhod James Jones

B. Noël Kivlin
Meyertons, Hood, Kivlin, Kowert & Goetzel
P.O. Box 398
Austin, TX 78767-0398

COMPUTER ASSEMBLY

BACKGROUND OF THE INVENTION

5

This invention relates to computer assemblies, and especially to computer assemblies that are employed as servers in data networks.

10

The assemblies may for instance be employed as servers for example in local area networks (LANs) or in wide area networks (WANs), telecommunications systems or other operations such as database management or as internet servers. Such servers may be used in so-called

15 "horizontally scaled" applications in which tens or hundreds of corresponding servers are employed as part of a distributed system.

20

A typical computer employed for such purposes will comprise a pair of processors mounted on a motherboard, together with power supply units (PSUs), and other components such as hard disc drives (HDDs), fans, digital video disc (DVD) players, memory modules etc. One or more of the processors, the host processor(s), provides

25 the main functions of the server, and may communicate

with a number of peripheral components, including communication ports, optionally via peripheral component interconnect (PCI) bridges in order to provide server operation.

5

In addition to the host processor(s), the assembly may include another processor, called the service processor or the remote management controller (RMC), which provides management functions for the system assembly. Such functions may include environmental monitoring, temperature monitoring of the enclosure, fan speed control, monitoring of various components of the assembly such as DVDs, memory cards, PSUs and the like.

10

15

Such assemblies are often provided as rack-mounted systems, and are housed in data centres which contain a large number of rows of assemblies extending for scores or even up to a hundred metres. When one of the assemblies malfunctions, a service engineer is called out to repair the assembly, normally by quiescing the assembly, replacing the component if it is in the form of a field-replaceable unit (FRU), and then testing the component. The assemblies typically have a console interface that communicates with the service processor, and, in use, is connected to a console in the form of a

20

25

personal computer located in the data centre so that the service engineer can log on to the console and obtain access to the relevant assembly. This operation can, however, take a significant amount of time. The service engineer will typically go to the remote console to diagnose the problem, walk to the assembly to repair the malfunction, and return to the console to verify the repair. For many service procedures this may involve repeating this operation a number of times, requiring the service engineer to walk hundreds of metres between the electronics assembly and the console. It would be possible for the service engineer to reduce the time taken for such a repair for example by taking a laptop computer with him to the relevant assembly for diagnostic purposes rather than returning to the console. However, the physical layout of such data centres is such that it is not easy or convenient to use a laptop at the racks of assemblies. Furthermore, operators of the data centres are understandably reluctant to allow individuals to bring computers into the building if they are then going to operate on the servers, in view of the sensitivity of the data held by the network and the possible damage to the network caused either by unauthorized reconfiguration of the network or by the introduction of a virus. Accordingly, it is desired that

all servicing or repair operations be conducted employing only computer equipment that is owned by the network operators.

5 SUMMARY OF THE INVENTION

According to one aspect of the present invention, there is provided a computer assembly which comprises:

- 10 (i) a housing;
- (ii) a host processor located within the housing;
- (iii) a service processor located within the housing for
15 providing system management functions within the computer assembly;
- (iv) a display that is located on the housing for displaying the status of components of the assembly
20 obtained from the service processor; and
- (v) one or more manual switches located on the housing for enabling a user to vary information displayed by the display and/or to alter the status of at
25 least one of the components.

Thus, according to this aspect of the invention, the service engineer can obtain diagnostic information on the assembly, and conduct testing of the assembly either
5 before or after repair thereof as required simply by actuating the switches to read the relevant data from the display. It is thus possible for the service engineer to effect repair of the assembly without the need to walk to the console or at least to minimise the number of
10 times it is necessary to walk to it.

At least in the broadest aspect of the invention, it is possible for all servicing operations to be effected at any time by accessing the electronics assembly. However,
15 according to another aspect, the invention provides a computer assembly which comprises:

- (i) a housing,
- 20 (ii) a host processor located within the housing;
- (iii) a service processor for providing system management functions within the computer system;
- 25 (iv) a display that is located on the housing for

displaying the status of components of the assembly
obtained from the service processor;

(v) one or more manual switches located on the housing
for enabling a user to vary information displayed
by the display and/or to alter the status of at
least one of the components; and

(vi) a console interface that communicates with the
service processor to enable system management
functions of the assembly to be monitored and/or
the status thereof to be modified from a console
connected thereto.

According to this aspect of the invention, the display
and/or switches on the assembly housing can be enabled
and/or disabled by means of signals sent thereto from the
console interface, thereby enabling the console to
determine which operations may be effected by the service
engineer. For example, the service engineer may be
allowed by the console to conduct routine service
operations or to replace any component that has
malfunctioned, but may not be allowed to conduct any
other operations such as re-configuring the assembly.

Thus, according to another aspect, the invention provides a computer system, which comprises a plurality of computer assemblies, each of which comprises:

- 5 (i) a housing;
-
- (ii) a host processor located within the housing;
-
- (iii) a service processor located within the housing for
10 providing system management functions within the
 computer assembly;
-
- (iv) a display that is located on the housing for
 displaying the status of components of the assembly
15 obtained from the service processor; and
-
- (v) one or more manual switches located on the housing
 for enabling a user to vary information displayed
 by the display and/or to alter the status of at
20 least one of the components;

the system including a console that can communicate with each of the assemblies and which can enable or disable the display and/or switches on any assembly either
25 completely or in part.

Also, according to yet another aspect, the invention provides a method of operating a computer system comprising a plurality of computer assemblies, each of which comprises:

5

(i) a housing;

(ii) a host processor located within the housing;

10

(iii) a service processor located within the housing for providing system management functions within the computer assembly;

(iv) a housing that encloses the computer assembly;

15

(v) a display that is located on the housing for displaying the status of components of the assembly obtained from the service processor;

20

(vi) one or more manual switches located on the housing for enabling a user to vary information displayed by the display and/or to alter the status of at least one of the components; and

25

(vi) a console that can communicate with each of the

assemblies and which can enable or disable the display and/or switches on any assembly either completely or in part;

5 which method comprises enabling the display and/or switches of an assembly that includes a component that has malfunctioned in order to allow the component to be repaired or replaced and to be tested, and then disabling the display and/or switches when the repair or
10 replacement has been effected.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be
15 described in detail by way of example with reference to the accompanying drawings, in which corresponding parts are given like reference numbers. In the drawings:

Figure 1 is a schematic block diagram showing part of
20 a network employing a number of electronics assemblies according to one embodiment of the present invention;

Figure 2 is a schematic view of a number of assemblies
25 forming a part of the network of figure 1;

Figure 3 is a schematic physical plan view of one form of computer assembly according to an embodiment of the present invention.

5 Figure 4 is a schematic block diagram showing the system architecture of an electronics assembly according to an embodiment of the invention;

10 Figure 5 is a detail taken from figure 2 showing components of the assembly housed in the front bezel thereof; and

Figure 6 is a schematic view of the front of one of the assemblies employed in the invention

15

DESCRIPTION OF PARTICULAR EMBODIMENTS

Referring now to the drawings, in which like reference numerals are used to designate corresponding elements, figure 1 shows part of a network in which one or more electronics assemblies according to one embodiment of the present invention may be employed. The network may for example be a local area network, or a wide area network and may be connected to other networks such as the internet. The network may comprise a firewall 1 for connection to the internet, for protecting the network

20

25

from viruses and unauthorized access, and for load
balancing of the network, switches 2 and 4 (there may be
many such switches in any one network) and customer
servers 6, 8 and 10. Thus, part of the network
5 comprising server 6, part of switch 2 and part of the
firewall 1 will form a network of customer A, while other
parts of the network will form networks of customers B
and C. All these networks will typically be housed in
one data centre that is run by a separate management
10 organisation 12 that will be able to access all parts of
the network. Clearly, even though the networks of the
different customers are physically located together, it
is important that data flow is contained within the
separate networks and is not allowed to cross between
15 them inadvertently. The management organisation 12 will
be in control of configuring and maintaining the network
among other things, and it will normally be relatively
insecure, relying on physical security of the building
and systems in order to prevent inappropriate transfer
20 of data.

Figure 2 shows a number of electronics assemblies 14
according to this embodiment of the invention that
provide servers for customers A, B and C in the network.
25 The assemblies 14 may be stacked together in racks, for

example nineteen inch, 23 inch or metric racks (only four are shown, but hundreds of such assemblies may be located in racks in a data centre) and may be connected to the switches 2 and 4, and to other parts of the network including a console 16 for controlling the management functions of the assemblies, that is to say, the configuration of the assemblies within the network, and managing the system functions of the assemblies.

Figure 3 shows a physical plan view of a typical narrow form factor computer that may provide a rack mounted server 14 in such a network.

The assembly may be designed to be a so-called high "RAS" system, that is to say, to have high reliability, availability and serviceability. As such, it is intended that the system will be operated with the minimum amount of down time, and indeed, such down time may be of the order of minutes per year.

The computer comprises an enclosure 21 that contains a motherboard 22 in the form of a printed circuit board (PCB) designed in a custom form-factor to fit the enclosure 21 and chosen to minimise the cabling withing the enclosure. The motherboard 22 carries the majority

of circuitry within the computer. On the motherboard are mounted one or (as shown here) two host processors or central processing units (CPUs) 24,26 each of which is provided with its own dedicated cooling in the form of an impingement fan that clips onto the CPU socket. Each processor 24, 26 may be provided with its own dedicated block of memory 27, 28 provided in the form of one or two banks of dual in-line memory modules (DIMMs) with a total of typically 256MB to 16GB block capacity.

A hardware cryptographic module (HCM) 30 may also be located on the motherboard. The HCM may be provided on a mezzanine card which plugs directly into the motherboard, and contains a co-processor providing cryptographic protocol acceleration support for security algorithms used in private community applications.

Two hard disc drives (HDDs) 32 and 34 are located at the front of the computer behind the front bezel 36. The drives are hot-pluggable and are accessible by removal of the bezel and EMI shield 38. Two internal HDDs plug directly into the motherboard via right-angled connectors located on the front edge of the motherboard 22.

Next to the HDDs may be arranged a system configuration

card reader (SCCR) 40 that is able to read a system configuration card (SCC) 42 inserted therein. The SCC may contain all relevant information concerning the computer, so that it is possible to replace one computer with another simply by inserting the original SCC into the new computer and replacing the hard disc drives with those of the original computer.

A removable media drive bay may be provided to allow optional fitting of a slimline (notebook style) digital video disc or digital versatile disc (DVD) drive 44 for reading CD and DVD media. The media transport loader is accessible through a slot in the enclosure bezel 36.

One or two 320W or 400W custom power supply units (PSUs) 46 may also be provided. In addition to the dedicated CPU fans, the assembly can be cooled by means of a row of fans 28 mounted between the motherboard and the media drive bays.

The computer may support input/output (I/O) expansion by means of peripheral component interconnect (PCI) cards that plug into expansion slots. These are accommodated by means of riser cards 49 that plug directly into the motherboard 22.

A number of I/O interfaces and sockets 50 are provided along the rear surface of the enclosure 1.

Figure 4 is a schematic representation of the system architecture of the computer system according to the invention.

Two host processors or CPUs 24 and 26 have an integer execution unit, a floating point and graphics unit, 32kB level 1 instruction cache, 64kB level 1 data cache, 1MB (256k x 32) level 2 data cache, a memory controller with error correction code (ECC) and a bus (JBUS) interface controller. Four DIMM sockets 27 and 28 are associated with each CPU.

The CPUs 24, 26 are connected in this configuration to two PCI bridges 60, 62 which provide interfaces to independent 64 bit PCI buses leading to various peripheral components such as the riser cards 49, HDDs 32 and 34, the HCM 30 etc. The PCI bridge 60 for CPU 24 is also connected to a PCI I/O device 64 referred to as "South Bridge". This is an integrated PCI sub system which provides an integrated drive electronics (IDE) controller, a universal serial bus (USB) controller, independent universal asynchronous receiver/transmitters

(UARTs), XBUS bridge and a power management controller. The PCI I/O device 64 also provides the console interface for enabling user access to the host processors 24 and 26.

5

A service processor or remote management controller (RMC) 70 is included for providing local and remote management services for control of the server. Such management service functions may include but are not limited to:

10

- 1) power management control;
- 2) environmental monitoring;
- 3) enclosure management and event logging;
- 4) fan control;
- 15 5) voltage rail monitoring;
- 6) component status monitoring.

20

The service processor is also responsible for monitoring and reporting the operational status of the system. The processor typically operates from the +5V standby rail and is capable of power cycling and resetting of the host system, and has dedicated flash ROM and synchronous dynamic RAM (SDRAM).

25

Peripheral devices that are required for the management

functions, may include, but are not limited to, the system configuration card reader (SCCR) 40, PCI clock generator 72, general purpose IO (GPIO) devices 74 field replaceable unit identification devices (FRUIDs) 76, a
5 "time-of-day" real time clock 77, and a system temperature monitor 78. These devices are provided on an inter-integrated circuit (I2C) management bus 80. In addition to the flash ROM and SDRAM, the service controller can access electrically erasable programmable
10 ROM (EEPROM) that is provided in the temperature monitor 78 via the I2C management bus 80.

As well as monitoring the environment and managing the peripheral devices, the service processor can communicate
15 with the PCI I/O device or console interface COM1 serial port 82. User access to the service processor 70 from the console 16 is available either through an ethernet port 84 NET_MGT, or through an asynchronous serial port
20 86 SERIAL/MGT. In this way, remote user access is available either to the service processor 70 for management functions, or to the host processor(s) 24 and 26 via the service processor 70.

In addition to the other components, a local management
25 centre (LMC) 90 may be connected to the I2C management

bus 80 for accessing information regarding the status of the assembly and for altering the status thereof. The local management centre may be located physically in the assembly housing 21, for example in the front bezel 36.

5

Figure 5 shows the LMC 90 in greater detail. The LMC 90 comprises for example a microcontroller 92 that is connected to the I2C management bus and is also connected to display, for example a liquid crystal display (LCD) 94 and a number of switches 96 in the form of a keypad. Figure 6 shows schematically how the LMC may be arranged in the front bezel 36 of the assembly housing.

10

15

20

The LMC 90 may have an alphanumeric LCD 94 which can be menu driven by means of "UP", "DOWN", "LEFT", "RIGHT" AND "ENTER" keys 97, 98, 99, 100 and 102 respectively, a power button 104 and four service indicator light emitting diodes (LEDs) 105, 106, 107 and 108. The LEDs 105 to 108 may be employed to indicate that the system is active, a service is required, the assembly is ready to be removed or any other state. The UP and DOWN keys may be employed to exit and enter various levels of the menu, while the LEFT and RIGHT keys enable scrolling of the menu items.

25

The alphanumeric LCD may be designed as a 16 column by 2 row display (although other sizes may be employed) which can operate in a number of modes:

5 1) Default Display Mode

This mode allows the display of system identification data. At the time of shipment, the default data should be the system serial number to allow ease of
10 identification in the event of a "dead-on-arrival" (DOA). The user may then have the option to set the default display to various identifiers such as: managed system hostname, Service processor network interface internet protocol (IP) address, managed system serial number,
15 managed system primary media access controller (MAC) address, or a customer defined ASCII string.

2) Service Event Default Mode

20 This mode allows the display to show that the system requires service attention, and may operate in conjunction with one of the LEDs, for example a "service required" LED. The display 94 may indicate the nature of the event that has occurred, for example a failure of
25 one of the components, and also the number of events.

3) Status Event Default Mode

This mode indicates the status of the assembly, for example when the operating system is running, booting, when it is undergoing self-testing or diagnostics or is on standby.

The LCD/keypad may enable an engineer to navigate through various menus in order to perform set tasks. Typical menus include:

1) Set Display Mode

This allows the user to select the default system identifier to be used in the default display mode.

2) View Event Log Mode

This menu allows the user to view previous service events and/or any changes in the status of the assembly. The log should contain all previous service events and or status changes.

3) Configure Mode

This menu shows and allows the user to enter parametric data, for example the configuration variables of the assembly, the hostname the internet protocol (IP) address subnet mask and the IP gateway of the assembly.

5

4) Test Mode

This menu allows the user to run a diagnostic test on the relevant component, e.g. a component that appears faulty or one that has just been replaced.

10

When Test Mode is selected, the up/down keys may be used to run through the list of available tests and the particular test may be selected by the enter key.

15

5) Exit

This menu exits the menu option and displays the default display.

20

Thus, as an example, when the service processor indicates that a service event has occurred, the display 94 will indicate the type of service event, and information identifying the event may be sent to the remote console 16 which will call out a service engineer. The engineer

25

can then proceed to the relevant assembly, view the event
and the service log on the display, perform any
diagnostic testing, replace the component if it is a
field-replaceable unit (FRU), repeat any diagnostic
5 testing, and move on to the next job.

Thus, it is possible to perform routine servicing and
maintenance without the need to connect or visit the
console. Indeed, if all the servers are implemented in
10 accordance with the invention the network administration
can remove the consoles to a locked office, possibly
remote from the datacentre, thereby improving security
and reducing costs.

15 In addition, by appropriate design of the menu, it is
possible to simplify the diagnosis of faults in the
assembly, thereby reducing the skill level of the
personnel required to service the assembly.

20 The LMC has the ability to allow a number of different
activities to be performed on the assembly, some of which
are non-invasive such as reading identification and
status of the assembly, while others are highly invasive
such as re-configuring the assembly or replacing and
25 testing FRUs, and it may not be appropriate for all

activities to be allowed all the time. According to one aspect of the present invention, the LMC 90 does not require authentication by password, and so, in view of this, the network system, comprising a number of such assemblies together with a console, may allow what can be described by authentication by proxy, that is to say, the system management may allow some, but not all, activities to be performed on the assembly by the LMC in some cases, and all activities to be performed in others, under authentication by the console 16 which will, typically, require access to be allowed using a password.

With such a system, the console 16 may communicate with each assembly and enable or disable the display 94 and keypad 96. Thus, it is possible for the system to operate under a number of options, for example:

Option 1 - All Enabled

With this option, the console enables all features of the LCD/keypad to allow testing and re-configuration of the assembly/network system. This will typically occur when the assembly is commissioned and when major changes are made to the network system.

Option 2 - All Disabled

With this option, no information can be obtained from the LCD, and access may be made to the assembly via the LMC 90.

Option 3 - Partially Enabled

This option may be employed during normal operation of the network system. This option could, for example, enable non-invasive activities such as allowing the LMC 90 to operate in display mode and also, optionally, to display the service log when required, but to disable re-configuration and testing of the assembly.

Other options may also be available. For example, the system may be arranged so that when the console 16 is notified of a service event (malfunction of a component) by the service processor 70, or that a service interval has expired the console 16 will automatically enable the test mode of the LCD/keypad to allow certain service operations to be performed such as testing power down etc, optionally after having notified the data management function of the service event. With this mode of operation, it is possible for the console 16 to change

the mode of operation automatically after completion of the service operation and testing, for example when the service engineer exits the menu, so that the LCD/keyboard is only partially enabled thereafter.

5

The scope of the present disclosure includes any novel feature or combination of features disclosed therein either explicitly or implicitly or any generalisation thereof irrespective of whether or not it relates to the claimed invention or mitigates any or all of the problems addressed by the present invention. The applicant hereby gives notice that new claims can be formulated to such features during prosecution of this application or of any such further application derived therefrom. In particular, with reference to the appended claims, features from dependent claims can be combined with those of the independent claims and features from respective independent claims can be combined in any appropriate manner and not merely in the specific combinations enumerated in the claims.

20